

EXHIBIT 4

From: Levy, Max I. <MLevy@mofo.com>
Sent: Monday, February 10, 2025 5:26 PM
To: Alex Frawley <AFrawley@susmangodfrey.com>
Cc: NYT-AI-SG-Service@simplelists.susmangodfrey.com; #NewYorkTimes-Microsoft-FDBR <NewYorkTimes-Microsoft-FDBR@faegredrinker.com>; NewYorkTimes_Microsoft_OHS@orrick.com; dailynews-ai-rfem@rothwellfigg.com; ai@loevy.com; KVP-OAI <KVPOAI@keker.com>; openaicopyrightlitigation.lwteam@lw.com; OpenAICopyright <OpenAICopyright@mofo.com>
Subject: [NYT] Responses to RFP Disputes

EXTERNAL Email

Hi Alex,

We've reviewed the previous emails regarding RFP disputes from The Times (attached) and have consolidated the correspondence into a single email. OpenAI responds below where the parties have ongoing disputes.

Ongoing Disputes

RFPs 67, 88, 99, 100: OpenAI sent an email to The Times on February 4 in response to its January 28 letter. We look forward to your response.

At the February 4 meet and confer, we were surprised to discover that the parties appear to have a disagreement as to what the Court instructed the parties to meet and confer about at the January 22 hearing. As we stated on the call, we understood the Court to (1) recognize that OpenAI has already agreed to provide The Times with various forms of financial information and related materials, (2) instruct The Times to review those materials and, (3) if after that review The Times believes those materials are insufficient, order the parties to meet and confer regarding any alleged deficiencies.

Contrary to the Court's instruction, The Times's position appears to be that the Court effectively (but not actually) ruled in The Times's favor that projections and valuations are relevant to copyright claims. Based on that position, during our call The Times maintained its position that it is entitled to additional financial information *regardless of whether other financial information OpenAI has agreed to produce would be sufficient to make the arguments that The Times says it needs the new information to support.*

The parties agreed to bring this dispute to the Court again to clarify the Court's prior guidance.

RFP 71: At the meet and confer, OpenAI confirmed its position that it believes the consolidated

financial statements to be sufficient given The Times's stated reasons for seeking product-specific financial information. Nevertheless, to resolve the dispute, OpenAI offered to produce relevant, non-privileged documents in its possession, custody, and control, if any, sufficient to show its revenue generated by product on an annual basis, to the extent that information is kept in the ordinary course of business. The Times clarified that it is seeking product-specific costs as well, suggesting that producing only product-specific revenues would not resolve the dispute. At the meet and confer, OpenAI requested The Times explain the relevance of product-specific costs. The Times agreed to follow-up with an explanation.

OpenAI also told The Times that it had not yet determined whether cost information is recorded by product in the ordinary course of business. The Times responded that a confirmation that such information is not kept in the ordinary course would resolve the dispute.

RFP 103: In the spirit of meeting and conferring, OpenAI has engaged in a healthy back-and-forth with The Times on whether it is entitled to the number of GPUs used by OpenAI. And we reiterate here that if The Times presents a persuasive argument, we would consider it.

As things stand, however, the only argument that the Times has presented is that it needs to know the number of GPUs so that it can rebut a hypothetical argument from OpenAI that it would "oppos[e] destruction of the models on the ground that OpenAI exhausted significant financial resources in building the models, including costs associated with GPUs." But as we stated on our last call, even if OpenAI were to make that argument (and it does not concede that it will), OpenAI has already produced consolidated financial statements, which include the full scope of the "significant financial resources" invested to build models. For the sake of clarity, OpenAI's expenses related to GPUs are captured in those statements.

We therefore cannot fathom why the financial information OpenAI has produced, which includes *the full scope of financial resources dedicated to building OpenAI's models*, would be insufficient for The Times to rebut an argument that "OpenAI exhausted significant financial resources in building the models." Nor, apparently, can The Times, given that it was unable to provide an explanation either in previous correspondence or at the meet and confer.

If The Times continues to believe it is entitled to information relating to the number of GPUs used by OpenAI, please help us understand what we're missing.

Resolved Disputes

RFP 70: We confirm that by "Search" we mean what we had previously been referring to as SearchGPT. We also confirm that OpenAI is running the search terms identified in the 1/24 response pursuant to the parties' agreement. We understand this to resolve the dispute over RFP 70.

RFP 101: We confirm that the scope of OpenAI's response to RFP 101 will include ChatGPT, including relevant features within the ChatGPT product such as Browse and Search, the OpenAI API, the Custom GPTs "News Summarizer," Remove Paywall," "Webpilot," and "WebGPT," and Microsoft's Copilot (the "relevant products"). We have revised the proposal below accordingly and understand

this dispute to be resolved.

OpenAI will agree to conduct a reasonable search for non-privileged documents in its possession, custody, or control, if any, (i) discussing OpenAI's evaluation or assessment of the value of including Journalism content in the training datasets used to train the relevant models and/or in connection with the use of RAG in the relevant products, (ii) discussing any evaluation or assessment conducted, generated, or commissioned by OpenAI relating to the value of including Journalism content in the text training datasets used to train the relevant models and/or the ability to display Journalism content in output of the relevant products, and (iii) discussing why OpenAI chose to use some Times content for training as opposed to other content and/or rely on Times content in connection with the 's use of RAG in the relevant products.